arXiv:1605.07919v1 [stat.ME] 25 May 2016

Compression and Conditional Emulation of Climate Model Output

Joseph Guinness[1] and Dorit Hammerling[2]

[1]*North Carolina State University, Department of Statistics*
[2]*National Center for Atmospheric Research*

Abstract

Numerical climate model simulations run at high spatial and temporal resolutions generate massive quantities of data. As our computing capabilities continue to increase, storing all of the data is not sustainable, and thus it is important to develop methods for representing the full datasets by smaller compressed versions. We propose a statistical compression and decompression algorithm based on storing a set of summary statistics as well as a statistical model describing the conditional distribution of the full dataset given the summary statistics. The statistical model can be used to generate realizations representing the full dataset, along with characterizations of the uncertainties in the generated data. Thus, the methods are capable of both compression and conditional emulation of the climate models. Considerable attention is paid to accurately modeling the original dataset–one year of daily mean temperature data–particularly with regard to the inherent spatial nonstationarity in global fields, and to determining the statistics to be stored, so that the variation in the original data can be closely captured, while allowing for fast decompression and conditional emulation on modest computers.

Keywords: Spatial-Temporal Data, Gaussian Process, Half-Spectral, Nonstationary, SPDE

1 Introduction

Development of high performance computing facilities with ever-increasing power allows climate scientists to resolve and study small-scale climate and weather phenomena with high resolution model simulations. The resulting deluge of climate data is now a serious challenge for the advancement of climate science (Kunkel et al., 2014). As a matter of fact, climate model data is routinely discarded or not saved at desired resolutions to fit within storage constraints. The problem is especially pertinent for large multi model ensemble projects such as the Coupled Model Comparison Projects (CMIP) used for the Intergovernmental Panel on Climate Change (IPCC) reports. CMIP5 storage requirements were approximately 2.5 PB, and a five-fold increase is expected, with an anticipated storage requirement of more than 10 PB, for the upcoming CMIP6 ensemble (Paul et al., 2015). Similar issues arise with projects investigating internal variability, such as the Community Earth System Model (CESM) large

[1]E-mail: jsguinne@ncsu.edu
[2]E-mail: dorith@ucar.edu

ensemble project (CESM-LE) (Kay et al., 2015), where the initial 30 ensemble members alone generated data exceeding 300 TB.

Compression methods are classically categorized as lossfree and lossy, where the former implies an exact reconstruction of the data, and the latter implies some loss of information (Sayood, 2012). The popular *gzip* method, for example, is a lossfree method. While the exact preservation of information is an attractive feature, lossfree compression methods have proven ineffective with many scientific applications due to the random nature of the trailing digits of the floating-point data (e.g., Lindstrom and Isenburg (2006), Bicer et al. (2013)). The application of lossy data compression to climate data is still in its infancy, so to obtain a better understanding of its effects, Baker et al. (2014) investigated lossy compression in the context of ensemble variability and determined that a compression rate of 5:1 was achievable for yearly-averaged variables. The method that provided the best results was the *fpzip* algorithm (Lindstrom and Isenburg, 2006), which truncates a specified number of least significant bits during the conversion of floating-point values to integers. One drawback, however, of *fpzip*, along with all other deterministic lossy compression algorithms, is the lack of uncertainty measures provided with the compressed data.

Recently, there has been interest in fitting statistical models to climate model output for the purpose of building computationally efficient model emulators (Castruccio et al., 2014, 2013) and for use as a tool for compressing the output. Previous work has viewed the parameters in the statistical model as the compressed object (Castruccio and Genton, 2015; Castruccio and Guinness, 2015), and thus if the model is well-specified, it can be used to generate emulated model runs whose mean and covariance structure matches that of the original dataset. However, these methods are not designed to reconstruct individual model runs, which is of interest for compressing large ensemble projects that investigate internal variability in the climate models. Our work is focused simultaneously on accurately modeling the mean and covariance structure in the original data and on accurately reconstructing the original data. In our approach, a set of summary statistics serves as a lower dimensional version of the original dataset, while a statistical model is used to represent the conditional distribution of the data given the summary statistics. This hybrid approach of storing both a statistical model and a set of summary statistics is highly flexible in that one can select which and how many statistics to store in order to target the representation of particular features of the original data to a desired accuracy. Furthermore, storing summary statistics can act as a buffer against model misspecification; one can store portions of the data that are difficult to represent with a statistical model, for example non-Gaussian behavior or discontinuities along land/ocean boundaries. This setup implies that only what is *not* stored needs to be modeled accurately, conditional on the stored statistics.

We consider the compression and decompression of one year of daily mean temperatures on a $1^\circ \times 1.25^\circ$ latitude $\times$ longitude grid, totaling 20 million values. In order to accurately model the major features of the data, we introduce a nonstationary space-time covariance model on the sphere that allows the characteristics of the time series to vary across the globe. The model is an instance of so-called *half spectral* models (Stein, 2005), in which the time series models are expressed in the spectral domain, and the coherence among the multiple time series is captured with a spatial covariance function whose parameters depend on frequency.

This model specification allows us to overcome the computational burden of computing Gaussian process approximate likelihoods because maps of Fourier coefficients can be considered approximately uncorrelated across frequencies. Further, we employ a stochastic partial differential equation (SPDE) representation (Lindgren et al., 2011) in the model for the Fourier coefficient maps, which induces sparsity in the inverse of the covariance matrix, which can be exploited to efficiently factor the matrix. To our knowledge, this is the first time a half spectral model has been combined with a computationally efficient representation for the Fourier coefficients, which allows for the analysis of truly massive space-time datasets with Gaussian process models.

We consider Fourier coefficients as summary statistics. This choice is motivated both by the computational advantanges discussed above, and by the fact that for many locations, especially for the ocean pixels, the variation across time can be captured by a few low-frequency coefficients. Compression is performed by fitting the model using conditional likelihoods and storing the statistical model and a set of summary statistics. We introduce a method for automatically selecting the Fourier coefficients based on a greedy selection algorithm, and we propose a variant of this algorithm that can be distributed and is hence computationally efficient for high performance computing infrastructures typically available at climate data centers. Decompression is performed by computing conditional expectations and conditional simulations given the summary statistics and the fitted model. We demonstrate that decompression can be performed quickly on a modest computer, having an end user with a laptop in mind. Finally, we evaluate the ability of the conditional simulations to accurately reflect the statistical properties of the small-scale features in the original data, and we conclude with a discussion.

2 The Data

The Community Earth System Model (CESM) large ensemble project (CESM-LE) (Kay et al., 2015) provides publicly available data from 40 ensemble model runs spanning the time period 1920-2100. We consider compression and decompression of one year (2081) of daily mean temperature fields from one ensemble member. The temperature values are reported on a 190×288 latitude $\times$ longitude grid, giving $n = 54,720$ temperature values each day, and thus the entire year of daily data has 19.97 million observations.

Let $Y(x,t)$ denote the daily mean temperature (in Celsius) at spatial location x and day $t \in 1,\ldots,T = 365$. It is clear that the first- and second-order properties of the time series vary substantially around the globe. To see this, we plot in Figure 1 time series from pixels nearest Chicago, Mumbai, the southern Atlantic Ocean, and Ross Island in Antarctica. Visual inspection of the time series indicates that the means, the seasonal cycles, the variances, and the autocorrelations are quite different in the four series. Such characteristics of the time series can be studied by examining their Fourier coefficients. The discrete Fourier transform (DFT) over time at each pixel is given by

$$\mathcal{Y}(\omega_k; x) := \frac{1}{\sqrt{T}} \sum_{t=1}^{T} Y(x,t) \exp(-i\omega_k t),$$

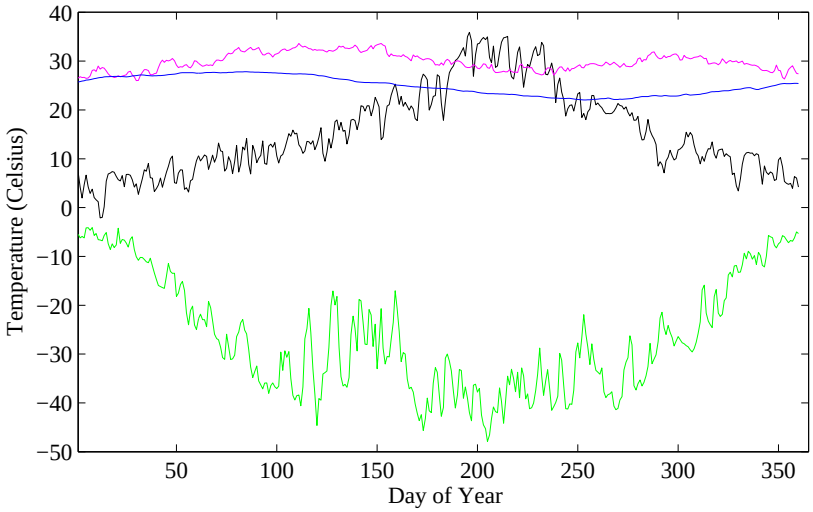

Figure 1: Temperature time series from Chicago (black), Mumbai (magenta), south Atlantic Ocean (blue), and Ross Island, Antarctica (green).

where $\omega_k = 2\pi k/T$ are the Fourier frequencies associated with T, and can be computed for all Fourier frequencies in $O(T \log T)$ floating point operations (flops) with fast Fourier transform (FFT) algorithms, and thus all of the DFTs at the $n = 190 \times 288 = 54,720$ spatial locations can be computed in $O(nT \log T)$ flops. To put it more concretely, all n DFTs are computed in a total of 0.54 seconds on a Macbook Air with a 1.7GHz Intel Core-i7 processor with 8GB RAM on Matlab R2013a. We refer to this computer henceforth as our "modest laptop."

Some of the major features of the spatial-temporal variation in the data can be summarized by making spatial maps of the Fourier coefficients $\mathcal{Y}(\omega_k, x)$, or functions thereof. In Figure 2(a), we plot a map of $T^{-1/2}\mathcal{Y}(\omega_0; x)$, which is simply the sample mean over time at each spatial location. In Figure 2(b), we plot the real part of $2T^{-1/2}\mathcal{Y}(\omega_1, x)$, which is variation in the data explained by $\cos(2\pi t/T)$, so it can be viewed as a signed strength of seasonal cycle–negative in the northern hemisphere, and positive in the southern hemisphere. In Figure 2(c), we plot

$$\widetilde{\sigma}(x) = \sqrt{\frac{1}{T-3}\sum_{k=2}^{T-2} |\mathcal{Y}(\omega_k; x)|^2},$$

which is equal to the sample standard deviation of the deseasonalized time series.

If $Y(x,t)$ is modeled as a stationary process in time with spectral density $f(\omega; x)$, we can use the Fourier coefficients to form a nonparametric estimate of the spectral densities of the time series, as in

$$\widetilde{f}(\omega_k; x) = \frac{1}{2\pi}\sum_{j=0}^{T-1} \alpha(k-j)|\mathcal{Y}(\omega_j; x)|^2,$$

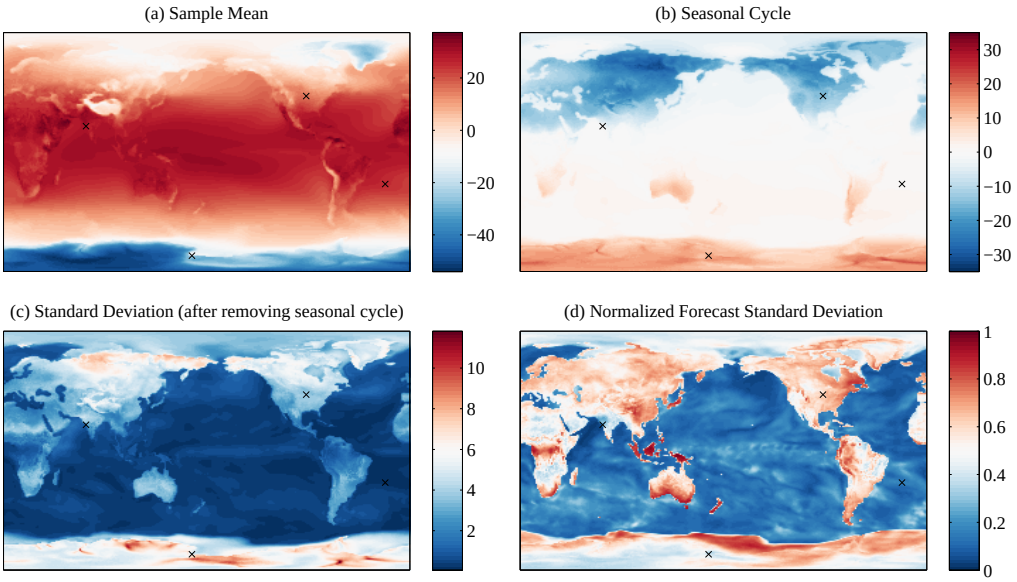

Figure 2: Maps of sample mean, seasonal cycle, standard deviation of deseasonalized data (all in Celsius), and normalized forecast standard deviation (unitless). The black crosses indicate the locations of the time series plotted in Figure 1.

where α is a smoothing kernel. In this exploratory data analysis, we use a Daniell window with bandwidth 17. Under the stationary-in-time assumption, Kolmogorov's formula (Brockwell and Davis, 2006, Theorem 5.8.1) gives the one-step-ahead prediction variance of $Y(x, t+1)$ based on an infinite past $Y(x,t), Y(x,t-1), \ldots$ in terms of the spectral density,

$$\mathrm{Var}(Y(x,t+1)|Y(x,t), Y(x,t-1), \ldots) = 2\pi \exp\left(\frac{1}{2\pi}\int_0^{2\pi} \log f(\omega; x) d\omega\right).$$

In Figure 2(d), we plot a map of the estimated one-step-ahead prediction standard deviation normalized by the sample standard deviation,

$$\widetilde{\sigma}(x)^{-1} 2\pi \exp\left(\frac{1}{T}\sum_{k=2}^{T-2} \log \widetilde{f}(\omega_k; x)\right)^{1/2}.$$

Very smooth time series tend to have normalized one-step-ahead prediction variances near zero, and those with a value near 1 are likely to be well-modeled by white noise. Panels (c) and (d) indicate that the deseasonalized time series are generally more variable over land than over the oceans, and further, the time series are smoother over the oceans since their normalized forecast standard deviations are smaller as well, giving empirical evidence for the qualitative features seen in Figure 1.

In Figure 3(a-b) we plot the real and imaginary parts of $\mathcal{Y}(\omega_{10}; x)/\widetilde{f}(\omega_{10}; x)^{1/2}$, a map of low-frequency Fourier coefficients normalized by their estimated spectral densities, and

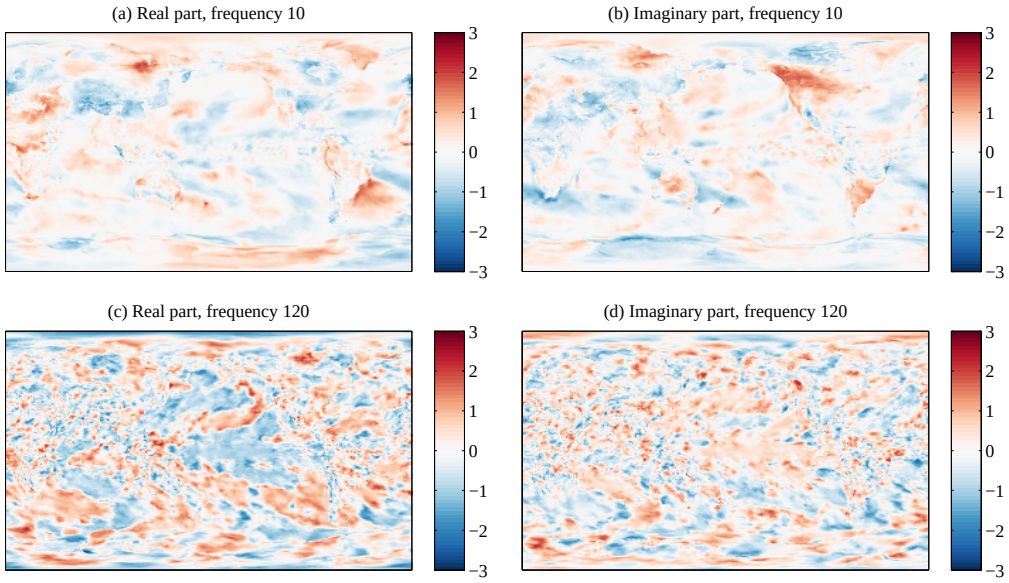

Figure 3: Maps of the real and imaginary parts of $\mathcal{Y}(\omega_k; x)/\sqrt{\widetilde{f}(\omega_k; x)}$, for $k = 10$ and $k = 120$.

in panels (c-d) we plot the real and imaginary parts of $\mathcal{Y}(\omega_{120}; x)/\widetilde{f}(\omega_{120}; x)^{1/2}$, a map of normalized high-frequency Fourier coefficients. We see that both the real and imaginary parts are each spatially correlated, and the spatial correlation is stronger at the lower frequency than it is at the higher frequency. This is consistent with the findings for wind data in Stein (2005), atmospheric pressure data in Stein (2009), and temperature data in Guinness and Stein (2013) and is also intuitive; low frequency fluctuations in temperature usually affect large regions simulataneously, whereas high frequency fluctuations can occur locally. The maps of normalized Fourier coefficients do not appear to deviate substantially from an assumption of isotropy. Possible exceptions include longer correlation ranges over the oceans and lack of dependence across land/ocean boundaries, especially at the lower frequency. In Section 3, we present a way to model the spatial maps of scaled Fourier coefficients with spatial processes. The spatial model allows us to predict the full maps in Figure 3 from a subset of the pixels via computationally efficient spatial interpolation.

3 Nonstationary Nonseparable Space-Time Model

Let $\mathbb{S}^2$ be a sphere, and $\mathbb{Z}$ be the integers. We model temperature at location $x \in \mathbb{S}^2$ and time $t \in \mathbb{Z}$ as

$$Y(x,t) = \mu(x,t) + \frac{1}{\sqrt{2\pi}} \int_0^{2\pi} \exp(i\omega t) d\widetilde{\mathcal{Y}}(\omega, x), \tag{1}$$

where $\widetilde{\mathcal{Y}}(\omega, x)$ is an orthogonal increment process in ω for every x, with

$$E(d\widetilde{\mathcal{Y}}(\omega, x)d\widetilde{\mathcal{Y}}(\omega, y)^*) = \sqrt{f(\omega; x)f(\omega; y)}C(x, y; \omega)d\omega,$$

where $*$ is complex conjugation, and $C(x, y; \omega)$ is a complex correlation function and thus $f(\omega; x)$ is the spectral density of the time series at location x. We assume $\mu(x, t)$ has the form

$$\mu(x, t) = \mu_0 + \mu_1 \exp(i\omega_1 t) + \mu_1^* \exp(i\omega_{T-1} t),$$

where μ_0 is real and μ_1 is complex.

Approximating the integral in (1) with a sum over T Fourier frequencies,

$$Y(x, t) \approx \mu(x, t) + \frac{1}{\sqrt{T}} \sum_{k=0}^{T-1} \exp(i\omega_k t)\mathcal{Y}(\omega_k, x),$$

along with the assumption that

$$E(\mathcal{Y}(\omega_k, x)\mathcal{Y}(\omega_j, y)^*) = \begin{cases} \sqrt{f(\omega_k; x)f(\omega_k; y)}C(x, y; \omega_k) & \text{if } k = j \\ 0 & \text{if } k \neq j, \end{cases}$$

leads to the usual approximation in the Whittle likelihood for multivariate time series (Whittle, 1953) and affords large computational savings because the maps of Fourier coefficients can be considered uncorrelated when computing likelihoods and simulations. Computational details are given in Subsections 4.1 and 4.3.

The model in (1) is a spatially nonstationary extension of the stationary model introduced in Stein (2005), which has been termed a "half-spectral" space-time model (Horrell and Stein, 2015) due to the fact that the temporal covariances are expressed in the frequency domain. The principal advantage of modeling space-time data in this way is that the spatial maps of Fourier coefficients $\mathcal{Y}(\omega_k, x)$ are approximately uncorrelated across ω_k. This brings enormous computational savings and allows us to flexibily specify and fit individual models to the spatial maps of Fourier coefficients at each frequency ω_k, which allows for space-time nonseparability. Half spectral models have been applied to high-frequency space-time atmospheric pressure data (Stein, 2009) and were generalized to handle high-frequency nonstationary in time temperature data (Guinness and Stein, 2013) with evolutionary spectra. Here, we allow the spectral densities $f(\omega, x)$ to vary with spatial location x, which gives nonstationary processes in space.

The inclusion of nonstationarity in space is motivated by the time series plots in Figure 1, where it can be observed that the temporal characteristics of the climate model data vary substantially across space. In particular, our model allows the spectral densities to depend on spatial location, so that the temporal behavior is stationary in time at any given spatial location, but the temporal behavior is nonstationary in space. In order to specify a flexible set of time series models that are not overwhelmingly burdensome to store, we propose a semiparametric modeling approach. We express the spectral density for the time series at location x as

$$f(\omega; x) = \exp(u_0(\omega) + \theta_1(x)u_1(\omega) + \theta_2(x)u_2(\omega) + \theta_3(x)u_3(\omega)),$$

where $\boldsymbol{\theta}(x) = (\theta_1(x), \theta_2(x), \theta_3(x))$ are unknown parameters associated with each spatial location x, and $\boldsymbol{u}(\omega) = (u_0(\omega), u_1(\omega), u_2(\omega), u_3(\omega))$ are data-dependent components of the log spectra. The first component $u_0(\omega)$ is the log average periodogram, averaged over pixels,

$$u_0(\omega) = \log\left(\frac{1}{n}\sum_{i=1}^{n}|\mathcal{Y}(\omega; x_i)|^2\right). \tag{2}$$

The three components u_1, u_2, and u_3 are the first three principal components of the log of smoothed and normalized periodograms. Details of the normalization and smoothing are given in Appendix A. Estimation details for $\boldsymbol{\theta}(x)$ are given in Section 4.

We model the coherences as

$$C(x, y; \omega) = K(\|x - y\|; \kappa(\omega)),$$

where K is the Matérn covariance function with variance and smoothness equal to 1, $\|\cdot\|$ denotes Euclidean distance in $\mathbb{R}^3$, and $\kappa(\omega)$ is an inverse range parameter that depends on frequency ω. The variance parameter is set to 1 because the overall variance of the process is absorbed into the spectral densities. The choice of smoothness equal to 1 is motivated by balancing fit to the data and computational feasibility. The maps of Fourier coefficients are not particularly smooth in space, so a small smoothness parameter is justified. An integer smoothness allows us to employ the stochastic partial differnetial equation (SPDE) representation for Matérn fields introduced in Lindgren et al. (2011), which induces sparsity in the inverse of the covariance matrix, and thus has computational benefits both for maximum likelihood estimation of $\kappa(\omega)$, but perhaps more importantly for allowing users to perform fast conditional simulations on modest computers, one of the central goals of this work. To our knowledge, this is the first instance in which the half-spectral model has been paired with a flexible and computationally efficient specification of the coherence functions, a pairing that allows for the analysis of truly massive space-time datasets with Gaussian process models.

The burden of storing the model parameters is: 3 numbers to represent μ_0 and μ_1, $3n$ numbers to represent $\boldsymbol{\theta}(x)$, $4T/2$ numbers to represent $\boldsymbol{u}(\omega)$, and $T/2$ numbers to represent $\kappa(\omega)$, and thus the total model storage burden is $3 + 3n + 5T/2$, which is roughly $(0.0083)nT$, or 0.83% of the size of the original dataset.

4 Model Estimation and Summary Statistic Selection

We consider the Fourier coefficients $\mathcal{Y}(\omega_k; x)$ as the summary statistics to store. The reason for this choice is two-fold. First, approximate independence across frequency of the Fourier coefficients greatly simplifies the conditional simulations given the stored Fourier coefficients–we can conditionally simulate the remaining Fourier coefficients at each frequency independently of those at other frequencies, and then transform back to the time-domain with an inverse discrete Fourier transform. Second, at many of the spatial locations, especially the ocean locations, most of the variance of the time series is concentrated on a few low-frequency Fourier

coefficients. As a consequence, many of the time series can be efficiently compressed by storing a small number of low-frequency Fourier coefficients. The total number of Fourier coefficients we can store is dictated primarily by the desired compression rate and the storage burden of the model. For example, at a 10:1 compression rate, we can store a total of 1.997 million numbers. Accounting for the model burden, this leaves us with 1.981 million numbers left to store. The total number of Fourier coefficients we can store is governed further by the ratio of zero frequency to non-zero frequency coefficients that are stored, since the non-zero frequency coefficients are complex, and thus count double. A further complicating matter is that, in addition to storing the Fourier coefficient $\mathcal{Y}(\omega_k; x)$, we must store k and x. We assume that each frequency-location pair (k, x) requires 8 bits on average, which we think is a conservative estimate and can be achieved with a standard delta encoding method (Sayood, 2012). Noting that the total number of stored Fourier coefficients is a somewhat complicated but easily computed function of the model, the compression rate, and the distribution of coefficients among frequencies, we denote the number of stored coefficients as N to simplify the discussion below.

Our task then is to select a set of N Fourier coefficients $\mathcal{Y}(\omega_1; x_1), \ldots, \mathcal{Y}(\omega_N, x_N)$ and model parameters $\boldsymbol{\theta}(x)$ and $\kappa(\omega)$ in order to best represent the original data and to accurately model the conditional distribution of the data given the Fourier coefficients. We consider the conditional likelihood to be a natural criterion for selecting the Fourier coefficients and the model parameters. The conditional likelihood is equal to the predictive density of the unstored coefficients, and thus larger conditional likelihoods correspond to sharper predictive distributions, since all of the densities are assumed to be Gaussian. Further, the joint density for the stored coefficients is irrelevant for this problem because there is no need to have an accurate statistical model for the data we store. The conditional likelihood given the stored coefficients is also a strictly proper scoring rule for prediction (Gneiting and Raftery, 2007).

While the conditional likelihood is a natural selection criterion, maximizing it over all choices of N Fourier coefficients and model parameters is an astronomically large combinatorial optimization problem (19.97 million choose 1.981 million is roughly 10 to the 2.92 millionth power). Thus, some severe restrictions on the search algorithm are necessary in order to make the selection computationally feasible. First, we estimate the time series model parameters once and for all by individually maximizing the Whittle likelihood for $\boldsymbol{\theta}(x)$ for each x. This part of the estimation takes on the order of 15 minutes on our modest laptop. We denote the estimates of the spectral densities as $\widehat{f}(\omega_k; x)$. The spatial coherence parameters $\kappa(\omega)$ are estimated initially using all of the data at each frequency. After this initial estimation, the Fourier coefficients are added sequentially in a greedy selection algorithm, to be described in Subsection 4.2. The selection of Fourier coefficients is periodically stopped to re-estimate each $\kappa(\omega)$ using the conditional loglikelihood given the current selection of coefficients, and then each $\kappa(\omega)$ is re-estimated at the end of the search to obtain the final estimates of the spatial coherence parameters. The lowest frequency Fourier coefficients determine the means and seasonal cycles, which are heavily influenced by the heterogeneous geography, and thus are expected to behave differently than high frequency coefficents. We mitigate this effect by fixing $\kappa(\omega_k) = 0.01$ for $k = 0, 1, 2$, which provided good performance in terms of the criteria discussed in Section 5. This subject is further discussed in Section 6.

4.1 Conditional Loglikelihoods

Before outlining the form of the conditional loglikelihoods, we first describe our notational convention. For variable $\mathcal{U}(\omega_k; x)$, which could be for example $\mathcal{Y}(\omega_k; x)$ or $\mathcal{Z}(\omega_k; x) = \mathcal{Y}(\omega_k; x)/\widehat{f}(\omega_k; x)^{1/2}$, define the vector

$$\boldsymbol{\mathcal{U}}(\omega_k) := \Big(\mathcal{U}(\omega_k; x_1), \ldots, \mathcal{U}(\omega_k; x_n)\Big).$$

We partition $\boldsymbol{\mathcal{U}}(\omega_k)$ as $(\boldsymbol{\mathcal{U}}_1(\omega_k), \boldsymbol{\mathcal{U}}_2(\omega_k))$ so that $\boldsymbol{\mathcal{U}}_1(\omega_k)$ is the vector of variables at frequency ω_k we store, and $\boldsymbol{\mathcal{U}}_2(\omega_k)$ is the vector of unstored variables at frequency ω_k. Further, define $\widehat{\boldsymbol{\mathcal{U}}}_2(\omega_k) = E(\boldsymbol{\mathcal{U}}_2(\omega_k)|\boldsymbol{\mathcal{U}}_1(\omega_k))$.

Under the model (1), each element of $\boldsymbol{\mathcal{Y}}(\omega_k)$ is approximately uncorrelated with each element of $\boldsymbol{\mathcal{Y}}(\omega_j)$ for $k \neq j$, and

$$\boldsymbol{\mathcal{Z}}(\omega_k) \sim CN(0, \Sigma(\kappa(\omega_k))).$$

We partition $\Sigma(\kappa(\omega_k))$ as

$$\Sigma(\kappa(\omega_k)) = \begin{bmatrix} \Sigma_{11}(\kappa(\omega_k)) & \Sigma_{12}(\kappa(\omega_k)) \\ \Sigma_{21}(\kappa(\omega_k)) & \Sigma_{22}(\kappa(\omega_k)) \end{bmatrix}, \tag{3}$$

where $\Sigma_{ij}(\kappa(\omega_k)) = E(\boldsymbol{\mathcal{Z}}_i(\omega_k)\boldsymbol{\mathcal{Z}}_j(\omega_k)^\dagger)$ and $\dagger$ is complex conjugate transpose.

Writing $Q(\kappa(\omega_k)) = \Sigma(\kappa(\omega_k))^{-1}$, and partitioning $Q(\kappa(\omega_k))$ similarly as in (3), the conditional expectation vector and covariance matrix of $\boldsymbol{\mathcal{Z}}_2(\omega_k)$ given $\boldsymbol{\mathcal{Z}}_1(\omega_k)$ are

$$E(\boldsymbol{\mathcal{Z}}_2(\omega_k)|\boldsymbol{\mathcal{Z}}_1(\omega_k)) = \widehat{\boldsymbol{\mathcal{Z}}}_2(\omega_k) = -Q_{22}(\kappa(\omega_k))^{-1}Q_{21}(\kappa(\omega_k))\boldsymbol{\mathcal{Z}}_1(\omega_k), \tag{4}$$

$$\mathrm{Var}(\boldsymbol{\mathcal{Z}}_2(\omega_k)|\boldsymbol{\mathcal{Z}}_1(\omega_k)) = Q_{22}(\kappa(\omega_k))^{-1}. \tag{5}$$

The conditional loglikelihood for $\kappa(\omega_k)$ is

$$CL_k(\kappa(\omega_k)) = -\frac{1}{2}\sum \log \widehat{f}(\omega_k; x) + \frac{1}{2}\log\det Q_{22}(\kappa(\omega_k)) - \frac{1}{2}(\boldsymbol{\mathcal{Z}}_2 - \widehat{\boldsymbol{\mathcal{Z}}}_2)^\dagger Q_{22}(\kappa(\omega_k))(\boldsymbol{\mathcal{Z}}_2 - \widehat{\boldsymbol{\mathcal{Z}}}_2),$$

where the first sum is over all locations x whose coefficients are not stored. Due to the approximately zero correlation between $\boldsymbol{\mathcal{Z}}(\omega_k)$ and $\boldsymbol{\mathcal{Z}}(\omega_j)$ for $j \neq k$, each $\kappa(\omega_k)$ can be estimated separately and in parallel by maximizing $CL_k(\kappa(\omega_k))$. The matrix $Q_{22}(\kappa(\omega_k))$ can be factored quickly due to the sparsity induced by the SPDE approximation. We factor the matrix with a Cholesky decomposition after permuting the rows and columns with a symmetric approximate minimum degree permutation (Amestoy et al., 1996), which encourages sparsity in the Cholesky factor.

4.2 Greedy Selection of Fourier Coefficients

Before selecting any Fourier coefficients, we estimate $\kappa(\omega_k)$ at each frequency using the marginal likelihood for all observations. Let $\widehat{\kappa}_0(\omega_k)$ denote the initial estimates. At this stage, we add selected coefficients before starting the search. Namely, for reasons discussed in Section 6, we

start the search with a low resolution grid of coefficients at frequencies ω_0 and ω_1. Define $\mathcal{R}(\omega_k; x) = \mathcal{Y}(\omega_k; x) - \widehat{\mathcal{Y}}(\omega_k; x)$, the residual for predicting Fourier coefficient $\mathcal{Y}(\omega_k; x)$ from the stored Fourier coefficients. Let

$$\mathcal{D}(\omega_k) = \max_x |\mathcal{R}(\omega_k; x)|^2.$$

Then we define $\omega_* = \arg\max \mathcal{D}(\omega_k)$, and select the M locations for which $|\mathcal{R}(\omega_*; x)|^2$ is largest, subject to the constraint that no two locations are within distance $d_{\min}$ of each other. We add the Fourier coefficients for the selected locations to $\boldsymbol{\mathcal{Y}}_1(\omega_*)$, and recompute $\mathcal{R}(\omega_*; x)$ and $D(\omega_*)$. We repeat this process until N Fourier coefficients have been added. This iterative procedure is stopped periodically to re-estimate $\kappa(\omega_k)$ using the conditional loglikelihoods given the current selection of stored Fourier coefficients. We denote the estimated coherence parameters at the jth stop as $\widehat{\kappa}_j(\omega_k)$ and the total number of stops with J.

If $M = 1$, this method corresponds to sequentially picking the coefficient with the largest residual under the current conditional model, and thus the choice of $d_{\min}$ would not be relevant. Since each iteration requires factoring $Q_{22}(\kappa(\omega_k))$, the algorithm is faster when M is larger because we add a large number of points at each iteration, and so fewer total iterations are required. The choice of $d_{\min}$ should in principle depend on M because if only a few locations are added per iteration (i.e. M is small), we should ensure that no two points are close to each other by making $d_{\min}$ large.

This particular method is intentionally iterative so that the estimates of $\kappa(\omega_k)$ are allowed to change when additional coefficients are stored, which is important since the maps of Fourier coefficients exhibit deviations from the assumption of isotropy. For example, since the coefficients are generally smoother over the oceans than over land, if we were to store all of the land coefficients at frequency ω_k, the conditional loglikelihood estimates $\widehat{\kappa}_J(\omega_k)$ would be smaller than the marginal loglikelihood estimates $\widehat{\kappa}_0(\omega_k)$. Additionally, storing coefficients for which $|\mathcal{R}(\omega_k; x)|^2$ is large encourages selection of coefficients that are either far in distance from already stored coefficients, or whose conditional mean is not accurately represented in the conditional distribution, a sign of possible anisotropy near x.

We have also developed an alternative version of the algorithm, where M locations are distributed to the frequencies ω_k proportional to $\mathcal{D}(\omega_k)$. The number of locations assigned to frequency ω_k, is defined as

$$m_k := \frac{\mathcal{D}(\omega_k)}{\sum_{k=0}^{T/2} \mathcal{D}(\omega_k)} M.$$

The remaining steps are identical to the sequential version, where, for each frequency ω_k, the m_k locations are chosen such that $|\mathcal{R}(\omega_k; x)|^2$ is largest subject to the minimum distance constraint and then $\mathcal{R}(\omega_k; x)$ and $\mathcal{D}(\omega_k)$ are recomputed. The periodic re-estimation of $\kappa(\omega_k)$ is also analogous. The main advantage of this "distributed" variant of the search algorithm is the ability to parallelize it, which can provide considerable savings in computational time. This is especially pertinent for lower compression ratios, for example 5:1, where a large number of locations are to be added.

We conducted a thorough analysis of the effect of various choices of M, J and $d_{\min}$ both on the timing of the algorithms and on the compression results. In the analysis presented

in Section 5 we present results for $M = 50$ (sequential) and $M = 7049$ (distributed), which corresponds to $m_k \approx 38$ per frequency on average, although the actual selection is heavily weighted towards lower frequencies. For both algorithms we used $J = 8$ and $d_{\min} = 0.05$. These settings provided decompressed datasets with small prediction errors and ran in a reasonable amount of time. This choice of $d_{\min}$ corresponds to a width of 2.3 pixels at the equator and guards against selecting too many pixels near the poles.

4.3 Decompression

Decompression is achieved by computing conditional expectations and conditional simulations given the stored coefficients. The conditional expectation of the scaled Fourier coefficients $\boldsymbol{\mathcal{Z}}_2(\omega_k)$ given $\boldsymbol{\mathcal{Z}}_1(\omega_k)$ can be computed as in Equation (4). Then the scaled Fourier coefficients are multiplied by the square root of the estimated spectral densities to obtain the Fourier coefficients, and finally, temperature values are obtained by computing inverse discrete Fourier transforms of each set of Fourier coefficients. The most demanding computational task is computing the conditional expectations, which involves solving a linear system with $Q_{22}(\kappa(\omega_k))$, accomplished first by computing the Cholesky factorization $LL^T = Q_{22}(\kappa(\omega_k))$, which is feasible due to the sparsity of $Q_{22}(\kappa(\omega_k))$, and then by solving the lower triangular sytem $Lv = -Q_{21}(\kappa(\omega_k))\boldsymbol{\mathcal{Z}}_1(\omega_k)$, followed by solving the upper triangular system $L^T u = v$, and setting $E(\boldsymbol{\mathcal{Z}}_2(\omega_k)|\boldsymbol{\mathcal{Z}}_1(\omega_k)) = u$. We use the Cholesky factor method because the Cholesky factor can also be used for simulating the conditional residuals e_2. For this task, we solve the upper triangular system $L^T e_2 = \varepsilon_2$, with $\varepsilon_2 \sim N(0, I)$, so that e_2 has covariance matrix $L^{-T}L^{-1} = Q_{22}^{-1}(\kappa(\omega_k))$.

The model has been chosen so that decompression can be performed quickly on a modest computer, while still allowing for substantial nonstationarity. Decompression can be parallelized as well by treating each frequency separately. The time required for decompression varies by strategy and the number of stored coefficients–storing more coefficients at frequency ω_k means that $Q_{22}(\kappa(\omega_k))$ is smaller and thus can be factored faster–but decompression generally took between 3 and 5 minutes without parallelization and between 1 and 2 minutes parallelizing over the 4 cores on our modest laptop.

5 Comparisons Between Original and Decompressed Datasets

We provide several criteria by which we evaluate the ability of the decompressed data to accurately represent the original data. Let $\widehat{Y}(x,t)$ denote the conditional expectation given the stored Fourier coefficients, and $\widetilde{Y}(x,t)$ denote a realization of the conditionally simulated data. The first performance criterion we consider is a pixelwise Root mean squared prediction error (RMSPE), defined as

$$s(x) = \sqrt{\frac{1}{T}\sum_{t=1}^{T}\Big(Y(x,t) - \widehat{Y}(x,t)\Big)^2}.$$

We plot maps of $s(x)$ for each level of compression in Figure 4. In Table 1, we report the square root of pixel area weighted averages of $s^2(x)$, averaged over water locations, land locations, and over all locations. As expected the errors decrease as the compression ratios decrease, and the average errors over ocean pixels are smaller than over land pixels. There are minor differences between the two algorithms, notably that the sequential version performs generally better over the ocean. The RMSPE differences are more prominent for the higher compression ratios, where the gains in computing times are also proportionally lower, so for practical purposes the choice of the selection algorithm could be a function of the compression needs and the available computing resources. For both algorithms, the computing times increase sublinearly with the number of coefficients selected. This finding is as expected, as we left the number of times $\kappa(\omega)$ is re-estimated, which is computationally demanding, constant for all compression ratios. The RMSPE values should be interpreted with reference to Figure 2, recalling that some of the deseasonalized time series in polar regions can be modeled as nearly white noise with standard deviations as large as 10 degrees Celsius, so a RMSPE of 1 degree Celsius is small in that context.

Selection algorithm	Comp. Ratio	land	ocean	all	runtime
sequential	20:1	0.7309	0.2901	0.4624	3.94
	10:1	0.4376	0.1933	0.2860	7.91
	5:1	0.2065	0.1171	0.1485	13.54
distributed	20:1	0.7355	0.3382	0.4871	1.16
	10:1	0.4368	0.2207	0.2994	1.57
	5:1	0.2042	0.1252	0.1522	2.40

Table 1: RMSPEs and run times for three different compression ratios and the two different greedy selection algorithms. Means are pixel area weighted averages. Runtimes are shown in hours.

In addition to producing small RMSPEs, the decompressed data also accurately capture the spatial-temporal dependence present in the original data. To see this, we plot spatial maps of averaged contrast variances, computed from the original data and from the conditionally simulated data. Define δ_{lat} and δ_{lon} to be displacement vectors of one pixel in latitude and longitude, respectively, so that $Y(x,t) - Y(x + \delta_{\text{lat}})$ is the contrast between the temperature at location x time t and the temperature at a location one pixel to the north at time t, and likewise $Y(x,t) - Y(x + \delta_{\text{lon}})$ is the contrast between the temperature at location x time t and the temperature at a location one pixel to the east at time t. In Figure 5, we plot maps

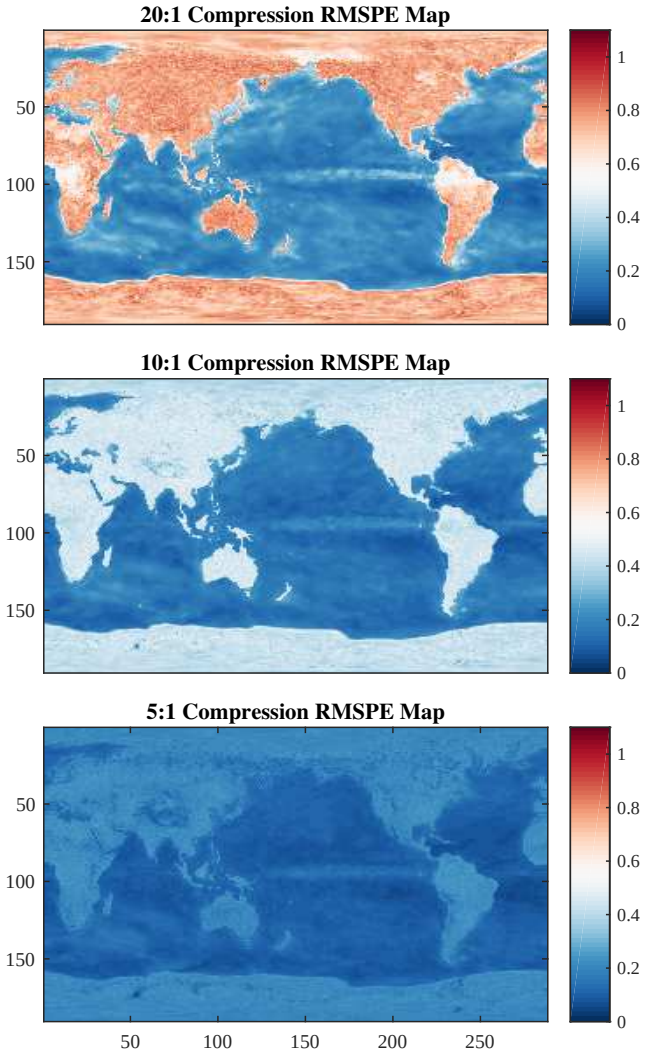

Figure 4: Maps of pixelwise RMSE. Units are degrees Celsius.

of the log of

$$\text{North-South}: \quad \frac{1}{T}\sum_{T=1}^{T}\Big(Y(x,t)-Y(x+\delta_{\text{lat}},t)\Big)^2,$$

$$\text{East-West}: \quad \frac{1}{T}\sum_{T=1}^{T}\Big(Y(x,t)-Y(x+\delta_{\text{lon}},t)\Big)^2,$$

$$\text{Temporal}: \quad \frac{1}{T-1}\sum_{T=1}^{T-1}\Big(Y(x,t)-Y(x,t+1)\Big)^2,$$

which are, respectively, the average north-south, east-west, and one step temporal contrast variances. Likewise, we plot the corresponding maps for a conditional simulation at each compression level. A remarkable feature of our methods, is that although we did not introduce any anisotropy in the spatial dependence, the contrast variances of the conditionally simulated data closely match those of the original data, with the accuracy increasing as the compression level decreases. The plotted maps are from the sequential selection algorithm; the distributed algorithm produced visually indistinguishable results.

In Figure 6, we reproduce the original time series plots from Figure 1, as well as the time series from the conditionally simulated data at the 10:1 compression level. The conditional simulations are able to reproduce the temporal dependence structures that vary substantially over space in the original data. The decompressed data also recover much of the small-scale structure as well. We point out here as well that even though the time series models were assumed to be stationary across time, the decompressed data accurately reflect some of the temporal nonstationarity as well; for example, in the Ross Island time series, the variance is larger in the middle of the year, and the decompressed data preserve this feature even though the assumed model is stationary across time.

6 Discussion

Model Assumptions: The model allows the spectral densities governing the temporal dependences to vary over the globe. Indeed, this is an obvious feature in the data, as seen in Figure 1. The model assumes that the standardized coefficients $\mathcal{Y}(\omega_k, x)/f(\omega_k, x)^{1/2}$ form an isotropic Gaussian process over the globe, and based on a visual inspection of Figure 3, a case could be made that this assumption is violated in that the correlation ranges appear to be longer over the oceans, and exhibit other nonstationary or non-Gaussian features, especially at low frequencies. However, there is a subtle point to be made: isotropy is only assumed in determining the conditional distributions for the unsaved coefficients given the saved ones. This is why it is important to use the conditinal loglikelihoods to estimate the spatial coherence parameters. No distributional assumptions are made about the saved coefficients–they are what they are, so we can use our choice of the saved coefficients to preserve nonstationarity or non-Gaussianity when it exists. Our algorithm for selecting the frequency-location pairs appears to be preserving nonstationarity automatically. Even though our model for the standardized Fourier

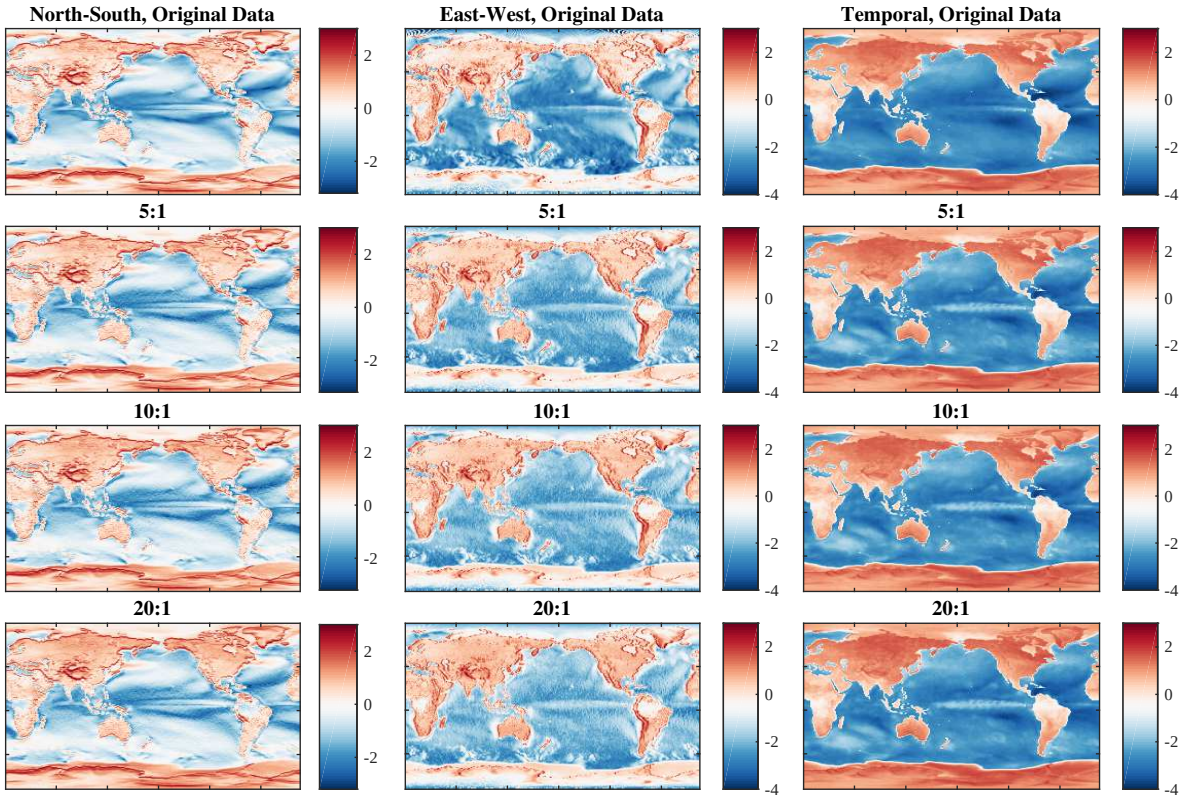

Figure 5: Maps of log contrast variances. First column are average North-South contrast variances, middle column are east-west contrast variances, and third column are one-step temporal contrast variances. First row is computed from the original data, second is from 5:1 decompressed data, third from 10:1 decompressed data, and the fourth row from 20:1 decompressed data.

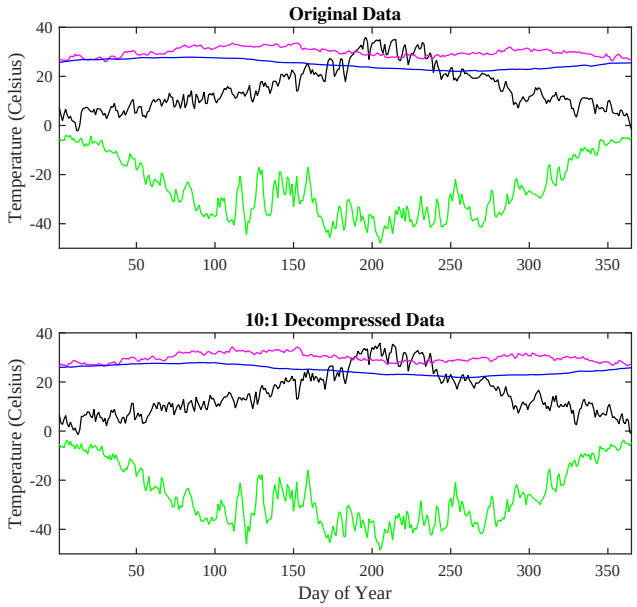

Figure 6: Original and 10:1 decompressed time series plots from Chicago (black), Mumbai (magenta), south Atlantic Ocean (blue), and Ross Island, Antarctica (green).

coefficients was isotropic in space, we can see in Figure 5 that the East-West and North-South contrast variances of the conditionally simulated data closely match the nonstationarity in the original data. We can see this effect in the time series plots as well. There is some evidence in Figures 1 and 6 that the time series are not stationary. For example, the Ross Island time series appears to have a larger variance in the winter months. Even though we modeled the time series as stationary processes, the conditionally simulated data in Figure 6 appear to be nonstationary as well.

Low Frequency Coefficients: At the lowest few frequencies, use of the conditional loglikelihood to estimate the spatial coherence parameters results in estimates of $\kappa(\omega)$ that correspond to processes that are visually rougher than the data. Since there is so much variation across the globe in the means and seasonal cycles, we hypothesize that departures from the Gaussian process assumption at those frequencies could be causing this behavior. A detailed investigation of non-Gaussian processes for the lowest frequencies is beyond the scope of this paper; however, we found that imposing some specific settings at the lowest frequencies improved the results. For $k = 0, 1$, we start the greedy search for coefficient selection by manually adding low resolution grids of 6,840 locations, which avoided selecting too many low frequency coefficients in early iterations of the distributed algorithm. For $k = 0, 1, 2$, we set $\widehat{\kappa}(\omega_k) = 0.01$, which improved the contrast variance maps.

Storing Numbers at a Fixed Precision: The original data are stored as single precision floating point numbers. In calculating the compression ratios, we have assumed throughout that all

parameters and Fourier coefficients are stored at the same precision, and that the frequency-location pairs are stored as an increasing sequence of integers to allow for delta encoding. Traditional compression algorithms make heavy use of storing different parts of the datasets at different precisions. We think that these ideas could certainly be applied to the present problem in conjunction with our proposed geostatistical compression methods, but we leave that problem for future work.

Polar Regions: The original data are reported on a regular latitude/longitude grid, which means that the original data devote a disproportionate amount of storage to the polar regions. An attractive feature of our methods is that the polar regions are treated in a natural fashion by the SPDE model, and our method for choosing the frequency-location pairs naturally avoids selecting too many locations near the poles. Further, the SPDE mesh can be easily adapted to non-regular grids in latitude and longitude.

Fixing Spectral Density Parameters: We did not attempt to perform a global optimization over the spatial range parameters and the spectral density parameters. Instead, we first fit the spectral density parameters individually at each pixel and then fixed them when maximizing the conditional loglikelihoods over the spatial range parameters. Performing the global optimization is not tractable because there are $3n$ time series parameters, and moreover, such a maximization could not be done separately at each frequency because the spectral density parameters affect the spectral densities of all frequencies. A different modeling strategy would be required for such global optimizations, possibly at the expense of having to assume a simpler, perhaps space-time separable, model.

Fixing the Compression Rate: The compression rate was fixed at three values, 20:1, 10:1, and 5:1, and then the model and coefficients were selected subject to the compression rate. This framework is chosen for simplicity of presentation, and the methods we propose could be easily adapted, for example, to a setting where one desires to reach a desired outcome, such as a specified root mean squared prediction error (RMSPE) or maximum absolute prediction error, with the best possible compression rate.

Code: We include Matlab code for reproducing the results in the online supplementary material.

Acknowledgements

We would like to acknowledge high-performance computing support from Yellowstone (ark:/85065/d7wd3xhc) provided by NCAR's Computational and Information Systems Laboratory, sponsored by the National Science Foundation. We further acknowledge support for long-term collaborative visits from the NSF Research Network on Statistics in the Atmosphere and Ocean Sciences (STATMOS) through grant DMS-1106862. This material is based upon work supported by the National Science Foundation under Grant No. 1406016. We would like to thank Allison Baker for sharing the data and helpful discussions and generous advice regarding its usage. We

would like to thank Anthony Tracy, Sophia Chen and William Kaufman for executing some of the HPC runs involved in this research.

A Principal Components of Smoothed Periodograms

The vector $\boldsymbol{u}_j = (u_j(\omega_0), \ldots, u_j(\omega_{T-1}))$ is defined as the jth principal component of

$$g(\omega_k, x) = \log\left(\sum_{\ell=0}^{T-1} \alpha(\ell - k)|\mathcal{Y}(\omega_\ell, x)|^2 / \exp(u_0(\omega_\ell))\right),$$

where $u_0(\omega_\ell)$ is defined in (2), and $\alpha(\ell - k)$ is a smoothing kernel defined as

$$\alpha(\ell) = \frac{1}{c}\exp(100(\cos(\omega_\ell) - 1)),$$

where c is a normalizing constant to ensure that α sums to 1.

References

Amestoy, P. R., Davis, T. A., and Duff, I. S. (1996). An approximate minimum degree ordering algorithm. *SIAM Journal on Matrix Analysis and Applications*, 17(4):886–905.

Baker, A., Xu, H., Dennis, J., Levy, M., Nychka, D., Mickelson, S., Edwards, J., Vertenstein, M., and Wegener, A. (2014). A methodology for evaluating the impact of data compression on climate simulation data. In *Proceedings of the 23rd International Symposium on High-performance Parallel and Distributed Computing*, HPDC '14, pages 203–214.

Bicer, T., Yin, J., Chiu, D., Agrawal, G., and Schuchardt, K. (2013). Integrating online compression to accelerate large-scale data analytics applications. *Parallel and Distributed Processing Symposium, International*, pages 1205–1216.

Brockwell, P. J. and Davis, R. A. (2006). *Time series: Theory and Methods*. Springer Science & Business Media.

Castruccio, S. and Genton, M. G. (2015). Compressing an ensemble with statistical models: an algorithm for global 3d spatio-temporal temperature. *Technometrics*, (just-accepted).

Castruccio, S. and Guinness, J. (2015). An evolutionary spectrum approach to incorporate large-scale geographical descriptors on global processes. *arXiv preprint arXiv:1507.03401*.

Castruccio, S., McInerney, D. J., Stein, M. L., Liu Crouch, F., Jacob, R. L., and Moyer, E. J. (2014). Statistical emulation of climate model projections based on precomputed GCM runs. *Journal of Climate*, 27(5):1829–1844.

Castruccio, S., Stein, M. L., et al. (2013). Global space–time models for climate ensembles. *The Annals of Applied Statistics*, 7(3):1593–1611.

Gneiting, T. and Raftery, A. E. (2007). Strictly proper scoring rules, prediction, and estimation. *Journal of the American Statistical Association*, 102(477):359–378.

Guinness, J. and Stein, M. L. (2013). Interpolation of nonstationary high frequency spatial–temporal temperature data. *The Annals of Applied Statistics*, 7(3):1684–1708.

Horrell, M. T. and Stein, M. L. (2015). Half-spectral space-time covariance models. *arXiv preprint arXiv:1505.01243*.

Kay, J., Deser, C., Phillips, A., Mai, A., Hannay, C., Strand, G., Arblaster, J., Bates, S., Danabasoglu, G., Edwards, J., Holland, M., Kushner, P., Lamarque, J.-F., Lawrence, D., Lindsay, K., Middleton, A., Munoz, E., Neale, R., Oleson, K., Polvani, L., and Vertenstein, M. (2015). The Community Earth System Model (CESM) large ensemble project: A community resource for studying climate change in the presence of internal climate variability. *Bulletin of the American Meteorological Society*, 96.

Kunkel, J. M., Kuhn, M., and Ludwig, T. (2014). Exascale storage systems–an analytical study of expenses. *Supercomputing Frontiers and Innovations*, 1(1).

Lindgren, F., Rue, H., and Lindström, J. (2011). An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 73(4):423–498.

Lindstrom, P. and Isenburg, M. (2006). Fast and efficient compression of floating-point data. *IEEE Transactions on Visualization and Computer Graphics*, 12:1245–1250.

Paul, K., Mickelson, S., Xu, H., Dennis, J. M., and Brown, D. (2015). Light-weight parallel Python tools for earth system modeling workflows. In *IEEE International Conference on Big Data*, pages 1985–1994.

Sayood, K. (2012). *Introduction to data compression*. Newnes.

Stein, M. L. (2005). Statistical methods for regular monitoring data. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 67(5):667–687.

Stein, M. L. (2009). Spatial interpolation of high-frequency monitoring data. *The Annals of Applied Statistics*, pages 272–291.

Whittle, P. (1953). The analysis of multiple stationary time series. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 125–139.